Proceedings of the 7th International Cryptology and Information Security Conference 2020 (CRYPTOLOGY2020)

No-Go Theorems for Data Privacy

Thomas Studer[*1]

[1]*Institute of Computer Science, University of Bern, Switzerland*

E-mail: thomas.studer@inf.unibe.ch
[*]*Corresponding author*

ABSTRACT

Controlled query evaluation (CQE) is an approach to guarantee data privacy for database and knowledge base systems. CQE-systems feature a censor function that may distort the answer to a query in order to hide sensitive information. We introduce a high-level formalization of controlled query evaluation and define several desirable properties of CQE-systems. Finally we establish two no-go theorems, which show that certain combinations of these properties cannot be obtained.

Keywords: Impossibility theorem, data privacy, controlled query evaluation, modal logic

1 INTRODUCTION

Controlled query evaluation (CQE) refers to a data privacy mechanism where the database (or knowledge base) is equipped with a censor function. This censor checks for each query whether the answer to the query would reveal sensitive information to a user. If this is the case, then the censor will distort the answer. Essentially, there are two possibilities how an answer may be distorted:

1. the CQE-system may refuse to answer the query (Sicherman et al., 1983) or
2. the CQE-system may give an incorrect answer, i.e. it lies (Bonatti et al., 1995).

This censor based approach has the advantage that the task of maintaining privacy is separated from the task of keeping the data. This gives more flexibility than an integrated approach (like hiding rows in a database) and guarantees than no information is leaked through otherwise unidentified inference channels. Controlled query evaluation has been applied to a variety of data models and control mechansims, see, e.g. Refs. Biskup (2000), Biskup and Bonatti (2001, 2004a,b), Biskup and Weibert (2008), Studer and Werner (2014).

No-go theorems are well-known in theoretical physics where they describe particular situations that are not physically possible. Often the term is used for results in quantum mechanics like Bell's theorem (Bell, 1964), the Kochen–Specker theorem (Kochen and Specker, 1968), or, for a more recent example, the Frauchiger–Renner paradox (Frauchiger and Renner, 2018). Nurgalieva and del Rio (2019) provide a modal logic analysis of the latter paradox. Arrow's theorem (Arrow, 1950) in social choice theory also is a no-go theorem stating that no voting system can be designed that meets certain given fairness conditions. Pacuit and Yang (2016) present a version of independence logic in which Arrow's theorem is derivable.

In the present paper we develop a highly abstract model for dynamic query evaluation systems like CQE. We formulate several desirable properties of CQE-systems in our framework and establish two no-go theorems saying that certain combinations of those properties are impossible. The main contribution of this paper is the presentation of the abstract logical framework as well as the high-level formulation of the no-go theorems. Note that some particular instances of our results have already been known (Biskup, 2000, Studer and Werner, 2014).

There are many different notions of privacy available in the literature. For our results, we rely on *provable privacy* (Stoffel and Studer, 2005, Stouppa and Studer, 2007), which is a rather weak notion of data privacy. Note that using a weak definition of privacy makes our impossibility theorems actually stronger since they state that under certain conditions not even this weak form of privacy can be achieved.

Clearly our work is also connected to the issues of lying and deception. Logics dealing with these notions are introduced and studied, e.g., by Ågotnes et al. (2018), Icard (2019), van Ditmarsch (2014).

In this version of the paper, we had to omit all proofs for lack of space. A version with full proofs is available in Ref. Studer (2020).

2 LOGICAL PRELIMINARIES

Let X be a set. We use $\mathcal{P}(X)$ to denote the power set of X. For sets Γ and Δ we use Γ, Δ for $\Gamma \cup \Delta$. Moreover, in such a context we write A for the singleton set $\{A\}$. Hence Γ, A stands for $\Gamma \cup \{A\}$.

Definition 2.1. *A* logic L *is given by*

1. *a set of formulas* Fml_L *and*
2. *a consequence relation* $\vdash_\mathsf{L}$ *for* L *that is a relation between sets of formulas and formulas, i.e.* $\vdash_\mathsf{L} \subseteq \mathcal{P}(\mathsf{Fml}_\mathsf{L}) \times \mathsf{Fml}_\mathsf{L}$ *satisfying for all* $A, C \in \mathsf{Fml}_\mathsf{L}$ *and* $\Gamma, \Delta \in \mathcal{P}(\mathsf{Fml}_\mathsf{L})$*:*
   (a) *reflexivity:* $\{A\} \vdash_\mathsf{L} A$*;*
   (b) *weakening:* $\Gamma \vdash_\mathsf{L} A \implies \Gamma, \Delta \vdash_\mathsf{L} A$*;*
   (c) *transitivity:* $\Gamma \vdash_\mathsf{L} C$ *and* $\Delta, C \vdash_\mathsf{L} A \implies \Gamma, \Delta \vdash_\mathsf{L} A$.

Transitivity is sometimes called *cut*. The previous definition gives us single conclusion consequence relations, which is sufficient for the purpose of this paper. For other notions of consequence relations see, e.g., Refs. Avron (1991) and Iemhoff (2016).

As usual, we write $\vdash_\mathsf{L} A$ for $\emptyset \vdash_\mathsf{L} A$. A formula A is called a *theorem of* L if $\vdash_\mathsf{L} A$.

We do not specify the logic L any further. The only thing we need is a consequence relation as given above. For instance, L may be classical propositional logic with $\vdash_{\mathsf{L}}$ being the usual derivation relation (see Section 4) or L may be a description logic with $\vdash_{\mathsf{L}}$ being its semantic consequence relation (Studer and Werner, 2014).

Definition 2.2.

1. *A logic* L *is called* consistent *if there exists a formula* $A \in \mathsf{Fml}_{\mathsf{L}}$ *such that* $\not\vdash_{\mathsf{L}} A$.
2. *A set* Γ *of* $\mathsf{Fml}_{\mathsf{L}}$*-formulas is called* L-consistent *if there exists a formula* $A \in \mathsf{Fml}_{\mathsf{L}}$ *such that* $\Gamma \not\vdash_{\mathsf{L}} A$.

We need a simple modal logic M over L.

Definition 2.3. *The set of formulas* $\mathsf{Fml}_{\mathsf{M}}$ *is given inductively by:*

1. *if* A *is a formula of* $\mathsf{Fml}_{\mathsf{L}}$, *then* $\Box A$ *is a formula of* $\mathsf{Fml}_{\mathsf{M}}$;
2. $\bot$ *is a formula of* $\mathsf{Fml}_{\mathsf{M}}$;
3. *if* A *and* B *are formulas of* $\mathsf{Fml}_{\mathsf{M}}$, *so is* $A \to B$, *too.*

As usual, the symbol $\bot$ denotes falsum and $\Box A$ means that A is known. We define the remaining classical connectives $\top$, $\wedge$, $\vee$, and $\neg$ in the standard way. Note that M is not a fully-fledged modal logic. For instance, it does not include nested modalities.

We give semantics to $\mathsf{Fml}_{\mathsf{M}}$-formulas as follows.

Definition 2.4. *An* M*-model* $\mathcal{M}$ *is a set of sets of* $\mathsf{Fml}_{\mathsf{L}}$*-formulas, that is*

$$\mathcal{M} \subseteq \mathcal{P}(\mathsf{Fml}_{\mathsf{L}}).$$

Definition 2.5. *Let* $\mathcal{M}$ *be an* M*-model. Truth of an* $\mathsf{Fml}_{\mathsf{M}}$*-formula in* $\mathcal{M}$ *is inductively defined by:*

1. $\mathcal{M} \Vdash \Box A$ *iff* $w \vdash_{\mathsf{L}} A$ *for all* $w \in \mathcal{M}$*;*

2. $\mathcal{M} \nVdash \bot$*;*

3. $\mathcal{M} \Vdash A \to B$ *iff* $\mathcal{M} \nVdash A$ *or* $\mathcal{M} \Vdash B$*.*

We use the following standard definition.

Definition 2.6. *Let* Γ *be a set of* $\mathsf{Fml}_{\mathsf{M}}$*-formulas.*

1. *We write* $\mathcal{M} \Vdash \Gamma$ *iff* $\mathcal{M} \Vdash A$ *for each* $A \in \Gamma$*.*

2. Γ *is called* satisfiable *iff there exists an* M*-model* $\mathcal{M}$ *with* $\mathcal{M} \Vdash \Gamma$*.*

3. Γ entails *a formula* A*, in symbols* $\Gamma \models A$*, iff for each model* $\mathcal{M}$ *we have that*
$$\mathcal{M} \Vdash \Gamma \quad \textit{implies} \quad \mathcal{M} \Vdash A.$$

3 PRIVACY

Definition 3.1. *A* privacy configuration *is a triple* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *that consists of:*

1. *the knowledge base* $\mathsf{KB} \subseteq \mathsf{Fml}_{\mathsf{L}}$*, which is only accessible via the censor;*

2. *the set of a priori knowledge* $\mathsf{AK} \subseteq \mathsf{Fml}_{\mathsf{M}}$*, which formalizes general background knowledge known to the attacker and the censor;*

3. *the set of secrets* $\mathsf{Sec} \subseteq \mathsf{Fml}_{\mathsf{L}}$*, which should be protected by the censor.*

A privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *satisfies the following conditions:*

1. KB *is* L*-consistent (consistency);*

2. $\{\mathsf{KB}\} \Vdash \mathsf{AK}$ *(truthful start);*

3. $\mathsf{AK} \not\models \Box s$ *for each* $s \in \mathsf{Sec}$ *(hidden secrets).*

Note that in the above definition, KB and Sec are sets of Fml_L-formulas while AK is a set of Fml_M-formulas. Thus AK may not only contain domain knowledge but also knowledge about the structure of KB. This is further explained in Section 4.

A *query* to a knowledge base KB is simply a formula of Fml_L.

Given a logic L, we can evaluate a query q over a knowledge base KB. There are two possible answers: t (true) and u (unknown).

Definition 3.2. *The evaluation function* eval *is defined by:*

$$\mathsf{eval}(\mathsf{KB}, q) := \begin{cases} t & \textit{if} \quad \mathsf{KB} \vdash_\mathsf{L} q \\ u & \textit{otherwise} \end{cases}$$

If the language of the logic L includes negation, then one may also consider an evaluation function that can return the value f (false), i.e. one defines $\mathsf{eval}(\mathsf{KB}, q) := f$ if $\mathsf{KB} \vdash_\mathsf{L} \neg q$. However, in the general setting of this paper, we cannot include this case.

A censor has to hide the secrets. In order to achieve this, it can not only answer t and u to a query but also r (refuse to answer). We denote the set of possible answers of a censor by

$$\mathbb{A} := \{t, u, r\}.$$

Let X be a set. Then X^ω denotes the set of infinite sequences of elements of X.

Definition 3.3. *A* censor *is a mapping that assigns an answering function*

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})} : \mathsf{Fml}_\mathsf{L}^\omega \longrightarrow \mathbb{A}^\omega$$

to each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$. *By abuse of notation, we also call the answering function* $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ *a* censor. *A sequence* $q \in \mathsf{Fml}_\mathsf{L}^\omega$ *is called* query sequence.

Usually, the privacy configuration will be clear from the context. In that case we simply use Cens instead of $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$.

Given a sequence s, we use s_i to denote its i-th element. That is for a query sequence $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$, we use q_i to denote the i-th query and $\mathsf{Cens}(q)_i$ to denote the i-th answer of the censor.

Example 3.1. *Let $A, B, C \in \mathsf{Fml}_{\mathsf{L}}$. We define a privacy configuration with* $\mathsf{KB} = \{A, C\}$, $\mathsf{AK} = \emptyset$, *and* $\mathsf{Sec} = \{C\}$. *A censor* Cens *yields an answering function* $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$, *which applied to a query sequence* $q = (A, B, C, \ldots)$ *yields a sequence of answers, e.g.,*

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q) = t, u, r, \ldots.$$

In this case, $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ *gives true answers since* $\mathsf{eval}(\mathsf{KB}, A) = t$ *and* $\mathsf{eval}(\mathsf{KB}, B) = u$ *and it protects the secret be refusing to answer the query C.*

Another option for the answering function would be to answer the third query with u, i.e., it would lie (instead of refuse to answer) in order to protect the secret.

A further option would be to always refuse the answer, i.e.

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q) = r, r, r, \ldots.$$

This, of course, would be a trivial (and useless) answering function that would, however, preserve all secrets.

In this paper, we will consider continuous censors only, which are given as follows.

Definition 3.4. *A censor* Cens *is* continuous *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and for all query sequences* $q, q' \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ *and all* $n \in \omega$ *we have that*

$$q|_n = q'|_n \quad \Longrightarrow \quad \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_n = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q')|_n$$

where for an infinite sequence $s = (s_1, s_2, \ldots)$, we use $s|_n$ to denote the initial segment of s of length n, i.e. $s|_n = (s_1, \ldots, s_n)$.

Continuity means that the answer of a censor to a query does not depend on future queries, see also Lemma 3.1.

A censor is called truthful if it does not lie.

Definition 3.5. *A censor* Cens *is called* truthful *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$*, all query sequences* $q = (q_1, q_2, \ldots)$*, and all sequences*

$$(a_1, a_2, \ldots) = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)$$

we have that for all $i \in \omega$

$$a_i = \mathsf{eval}(\mathsf{KB}, q_i) \quad \textit{or} \quad a_i = r.$$

Hence a truthful censor may refuse to answer a query in order to protect a secret but it will not give an incorrect answer.

In the modal logic M over L, we can express what knowledge one can gain from the answers of a censor to a query. This is called the content of the answer.

Definition 3.6. *Given an answer* $a \in \mathbb{A}$ *to a query* $q \in \mathsf{Fml}_\mathsf{L}$*, we define its* content *as follows:*

$$\begin{aligned} \mathsf{cont}(q, t) &:= \Box q \\ \mathsf{cont}(q, u) &:= \neg\Box q \\ \mathsf{cont}(q, r) &:= \top \end{aligned}$$

Assume that we are given a privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and a censor* Cens*. We define the content of the answers of the censor to a query sequence* $q \in \mathsf{Fml}_\mathsf{L}^\omega$ *up to* $n \in \omega$ *by*

$$\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) := \bigcup_{1 \le i \le n} \{\mathsf{cont}(q_i, a_i)\} \cup \mathsf{AK}$$

where $a = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)$*. Note that here we have also included the a priori knowledge.*

The following is a trivial observation showing the role of continuity.

Lemma 3.1. *Let* Cens *be a continuous censor. The content function is monotone in the second argument: for* $m \leq n$ *we have*

$$\mathsf{cont}(\mathsf{Cens}(q), m) \subseteq \mathsf{cont}(\mathsf{Cens}(q), n).$$

We call a censor credible if it does not return contradicting answers.

Definition 3.7. *A censor* Cens *is called* credible *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and for every query sequence q and every* $n \in \omega$*, the set* $\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n)$ *is satisfiable.*

Definition 3.8. *The* full content *of a knowledge base* KB *is given by*

$$\mathsf{full}(\mathsf{KB}) := \bigcup_{A \in \mathsf{Fml}_{\mathsf{L}}} \mathsf{cont}(A, \mathsf{eval}(\mathsf{KB}, A)).$$

Lemma 3.2. *For any knowledge base* KB*, we have that*

$$\{\mathsf{KB}\} \Vdash \mathsf{full}(\mathsf{KB}).$$

Lemma 3.3. *We let* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *be a privacy configuration. Further we let* $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ *be a truthful censor. For every query sequence q and* $n \in \omega$*, we have that*

$$\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) \subseteq \mathsf{full}(\mathsf{KB}) \cup \{\top\} \cup \mathsf{AK}.$$

The following corollary is a generalization of Cor. 30 in Ref. Studer and Werner (2014).

Corollary 3.1. *Every truthful censor is credible.*

There are several properties that a 'good' censor should fulfil. We call a censor effective if it protects all secrets.

Definition 3.9. *A censor* Cens *is called* effective *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and for every query sequence* $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ *and every* $n \in \omega$*, we have*

$$\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) \not\models \Box s \quad \textit{for each } s \in \mathsf{Sec}$$

A ‘good’ censor should only distort an answer to a query when it is absolutely necessary, i.e. when giving the correct answer would leak a secret. We call such a censor minimally invasive.

Definition 3.10. *Let* Cens *be an effective and credible censor. This censor is called* minimally invasive *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and for each query sequence* $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$*, we have that whenever*

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i \neq \mathsf{eval}(\mathsf{KB}, q_i),$$

replacing

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i \quad \textit{with} \quad \mathsf{eval}(\mathsf{KB}, q_i)$$

would lead to a violation of effectiveness or credibility, that is for any censor $\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ *such that*

$$\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_{i-1} = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_{i-1}$$

and

$$\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i = \mathsf{eval}(\mathsf{KB}, q_i)$$

we have that for some n

$$\mathsf{cont}(\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) \models \Box s \quad \textit{for some } s \in \mathsf{Sec}$$

or

$$\mathsf{cont}(\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) \textit{ is not satisfiable}.$$

It is a trivial observation that a truthful, effective and minimally invasive censor has to answer the same query always in the same way.

Lemma 3.4. *Let* Cens *be a truthful, effective and minimally invasive censor. Further let* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *be a privacy configuration and* q *be a query sequence with* $q_i = q_j$ *for some* i, j*. Then*

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_j.$$

Consider a truthful, effective, continuous and minimally invasive censor and a given query sequence. If the censor lies to answer some query, then giving the correct answer would immediately reveal a secret.

Lemma 3.5. *Let* Cens *be a truthful, effective, continuous and minimally invasive censor. Further let* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *be a privacy configuration and* q *be a query sequence. Let* i *be the least natural number such that*

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i \neq \mathsf{eval}(\mathsf{KB}, q_i).$$

Let $\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ *be such that*

$$\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_{i-1} = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_{i-1}$$

and

$$\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i = \mathsf{eval}(\mathsf{KB}, q_i).$$

Then it holds that

$$\mathsf{cont}(\mathsf{Cens}'_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), i) \models \Box s \quad \textit{for some } s \in \mathsf{Sec}.$$

Next we define the notion of a repudiating censor, which garantees that there is always a knowledge base in which no secret holds and which, given as input to the answering function, produces the same results as the actual knowledge base. Hence this definition provides a version of plausible deniability for all secrets.

Definition 3.11. *A censor* Cens *is called* repudiating *iff for each privacy configuration* $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ *and each query sequence* q*, there are knowledge bases* KB_i *(*$i \in \omega$*) such that*

1. $(\mathsf{KB}_i, \mathsf{AK}, \mathsf{Sec})$ *is a privacy configuration for each* $i \in \omega$*;*
2. $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_n = \mathsf{Cens}_{(\mathsf{KB}_n,\mathsf{AK},\mathsf{Sec})}|_n$*, for each* $n \in \omega$*;*
3. $\mathsf{KB}_i \not\models_{\mathsf{L}} s$ *for each* $s \in \mathsf{Sec}$ *and each* $i \in \omega$*.*

Now we can establish our first no-go theorem, which is a generalization of Th. 50 in Ref. Studer and Werner (2014).

Theorem 3.1 (First No-Go Theorem)**.** *A continuous and truthful censor satisfies at most two of the properties effectiveness, minimal invasion, and repudiation.*

4 NON-REFUSING CENSORS

In this section we study censors that do not refuse to answer a query.

Definition 4.1. *A censor is* non-refusing *if it never assigns the answer r to a query.*

Of course, a non-refusing censor has to lie in order to keep the secrets. That means if a censors of this kind shall be effective, then it cannot be truthful.

Even if we consider lying censors, we work with the assumption that

$$\text{an attacker believes every answer of the censor.} \tag{1}$$

Otherwise, we are in a situation where an attacker cannot believe any answer because the attacker does not know which answers are correct and which are wrong, which means that any answer could be a lie. In that case, querying a knowledge base would not make any sense at all.[1]

Because of the assumption (1), we can use our notions of effectiveness (Definition 3.9) and credibility (Definition 3.7) also in the context of lying censors: an attacker should not believe any secret and the beliefs should be satisfiable.

Theorem 3.1 about truthful censors did not make any assumptions on the underlying logic L. The next theorem about non-refusing censors is less general as it is based on classical logic. We will use $a, b, c, \ldots$ for atomic propositions and $A, B, C, \ldots$ for arbitrary formulas.

Moreover, we assume that the knowledge base KB only contains atomic facts (we say KB is *atomic*). That is if $F \in \mathsf{KB}$, then F is either of the form p or of the form $\neg p$ where p is an atomic proposition. Hence we find that if $\mathsf{KB} \vdash_{\mathsf{L}} a \to b$ for two distinct atomic propositions a and b, then $\mathsf{KB} \vdash_{\mathsf{L}} \neg a$ or $\mathsf{KB} \vdash_{\mathsf{L}} b$. We can formalize this using the set of a priori knowledge by letting

$$\Box(a \to b) \to (\Box\neg a \vee \Box b) \in \mathsf{AK}.$$

[1] This is, of course, not completely true. It is possible to distort knowledge bases in such a way that privacy is preserved but statistical inferences are still informative, see, e.g. Ref. du Pin Calmon and Fawaz (2012).

Now we can establish our second no-go theorem, which is a generalization of the results of Biskup (2000).

Theorem 4.1 (Second No-Go Theorem)**.** *Let* L *be based on classical logic. A continuous and non-refusing censor cannot be at the same time effective and minimally invasive.*

To avoid this problem, a censor must not only protect the single elements of Sec but also their disjunction (Biskup, 2000). Note that protecting the disjunction of all secrets is not as simple as it sounds. Consider, for instance, a hospital information system that should protect the disease a patient is diagnosed with. In this case, protecting the disjunction of all secrets means protecting the information that the patient has some disease. This, however, is not feasible as it is general background knowledge that everybody who is a patient in a hospital has some disease. Worse than that, sometimes the disjunction of all secrets may even be a logical tautology, which cannot be protected.

5 CONCLUSION

In this paper, we have established two no-go theorems for data privacy using tools from modal logic. We are confident that logical methods will play an important role for finding new impossibility theorems or for better understanding already known ones, see, e.g., the logical analyses carried out by Nurgalieva and del Rio (2019) and Pacuit and Yang (2016).

Another line of future research relates to the fact that refusing to answer a query can give away the information that there exists a secret that could be infered from some other answer. Similar phenomena may occur in multi-agent systems when one of the agents refuses to communicate. For example, imagine the situation of an oral exam where the examiner asks a question and the student keeps silent. In this case the examiner learns that the student does not know the answer to the question for otherwise the student would have answered.

It is also possible that refusing an answer can lead to knowing that someone else knows a certain fact. Consider the following scenario. A father enters a room where his daughter is playing and he notices that one of the toys is in pieces. So he asks who has broken the toy. The daughter does not want to betray her brother (who actually broke it) and she also does not want to lie. Therefore, she refuses to answer her father's question. Of course, then the father knows that his daughter knows who broke the toy for otherwise the daughter could have said that she does not know.

We believe that it is worthwhile to study the above situations using general communication protocols that include the possibility of refusing an answer and to investigate the implications of refusing in terms of higher-order knowledge.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers, who helped to improve this paper substantially. This work is supported by the Swiss National Science Foundation grant 200020_184625.

REFERENCES

Ågotnes, T., van Ditmarsch, H., and Wang, Y. (2018). True lies. *Synthese*, 195(10):4581–4615.

Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346.

Avron, A. (1991). Simple consequence relations. *Inf. Comput.*, 92(1):105–139.

Bell, J. S. (1964). On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1:195–200.

Biskup, J. (2000). For unknown secrecies refusal is better than lying. *Data and Knowledge Engineering*, 33(1):1–23.

Biskup, J. and Bonatti, P. A. (2001). Lying versus refusal for known potential secrets. *Data and Knowledge Engineering*, 38(2):199–222.

Biskup, J. and Bonatti, P. A. (2004a). Controlled query evaluation for enforcing confidentiality in complete information systems. *International Journal of Information Security*, 3(1):14–27.

Biskup, J. and Bonatti, P. A. (2004b). Controlled query evaluation for known policies by combining lying and refusal. *Annals of Mathematics and Artificial Intelligence*, 40(1):37–62.

Biskup, J. and Weibert, T. (2008). Keeping secrets in incomplete databases. *International Journal of Information Security*, 7(3):199–217.

Bonatti, P. A., Kraus, S., and Subrahmanian, V. S. (1995). Foundations of secure deductive databases. *Transactions on Knowledge and Data Engineering*, 7(3):406–422.

du Pin Calmon, F. and Fawaz, N. (2012). Privacy against statistical inference. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1401–1408. IEEE.

Frauchiger, D. and Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications*, 9.

Icard, B. (2019). *Lying, deception and strategic omission : definition et evaluation*. PhD thesis, Universit Paris Sciences et Lettres.

Iemhoff, R. (2016). Consequence relations and admissible rules. *Journal of Philosophical Logic*, 45(3):327–348.

Kochen, S. and Specker, E. (1968). The problem of hidden variables in quantum mechanics. *Indiana Univ. Math. J.*, 17:59–87.

Nurgalieva, N. and del Rio, L. (2019). Inadequacy of modal logic in quantum settings. In Selinger, P. and Chiribella, G., editors, *Proceedings 15th International Conference on Quantum Physics and Logic, QPL 2018, Halifax, Canada, 3-7th June 2018.*, volume 287 of *EPTCS*, pages 267–297.

Pacuit, E. and Yang, F. (2016). Dependence and independence in social choice: Arrow's theorem. In Abramsky, S., Kontinen, J., Väänänen, J., and Vollmer, H., editors, *Dependence Logic: Theory and Applications*, pages 235–260. Springer.

Sicherman, G. L., De Jonge, W., and Van de Riet, R. P. (1983). Answering queries without revealing secrets. *ACM Trans. Database Syst.*, 8(1):41–59.

Stoffel, K. and Studer, T. (2005). Provable data privacy. In Andersen, K. V., Debenham, J., and Wagner, R., editors, *Database and Expert Systems Applications*, pages 324–332. Springer.

Stouppa, P. and Studer, T. (2007). A formal model of data privacy. In Virbitskaite, I. and Voronkov, A., editors, *Perspectives of Systems Informatics*, pages 400–408. Springer.

Studer, T. (2020). No-go theorems for data privacy. E-print 2005.13811, arXiv.org.

Studer, T. and Werner, J. (2014). Censors for boolean description logic. *Transactions on Data Privacy*, 7:223–252.

van Ditmarsch, H. (2014). Dynamics of lying. *Synthese*, 191(5):745–777.